BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

As rescanning documents *will not* correct images, please do not report the images to the Image Problem Mailbox.

Results 61 - 80 of 200 Result page: previous 1 2 3 **4** 5 6
Best 200 shown

61 Difficult-path branch prediction using subordinate microthreads
Robert S. Chappell, Francis Tseng, Adi Yoaz, Yale N. Patt
May 2002 ACM SIGARCH Computer Architecture News , Proceedings of the 29th a architecture, Volume 30 Issue 2
Full text available: pdf(1.14 MB) Additional Information: full citation, abstra

Branch misprediction penalties continue to increase as microprocessor cores be prediction accuracy remains an important challenge. Simultaneous Subordinate improve branch prediction accuracy. SSMT machines run multiple, concurrent r We propose to dynamically construct microthreads that can speculatively and a frequently mis ...

62 The KScalar simulator
J. C. Moure, Dolores I. Rexachs, Emilio Luque
March 2002 Journal on Educational Resources in Computing (JERIC), Volume
Full text available: pdf(493.35 KB) Additional Information: full citation, abstract, ref

Modern processors increase their performance with complex microarchitectural difficult to understand and evaluate. KScalar is a graphical simulation tool that students to analyze the performance behavior of a wide range of processor mi scalar pipeline, to a detailed out-of-order, superscalar pipeline with non-blockir

Keywords: Education, pipelined processor simulator

63 Formalizing the safety of Java, the Java virtual machine, and Java card
Pieter H. Hartel, Luc Moreau
December 2001 ACM Computing Surveys (CSUR), Volume 33 Issue 4
Full text available: pdf(442.86 KB) Additional Information: full citation, abstract, referen

We review the existing literature on Java safety, emphasizing formal approach footprint devices such as smartcards. The conclusion is that although a lot of g is needed to build a coherent set of machine-readable formal models of the wh formidable task but we believe it is essential to build trust in Java safety, and t ...

Keywords: Common criteria, programming

64 Novel ideas: A design space evaluation of grid processor architectures
Ramadass Nagarajan, Karthikeyan Sankaralingam, Doug Burger, Stephen W. Ke
December 2001 Proceedings of the 34th annual ACM/IEEE international sympos
Full text available: pdf(1.29 MB) Additional Information: full citation, abstract, re

In this paper, we survey the design space of a new class of architectures callec architectures are designed to scale with technology, allowing faster clock rates superior instruction-level parallelism on traditional workloads and high perform consists of an array of ALUs, each with limited control, connected by a thin ope

65 Superscalar architectures: Reducing the complexity of the register file in c
Rajeev Balasubramonian, Sandhya Dwarkadas, David H. Albonesi
December 2001 Proceedings of the 34th annual ACM/IEEE international sympos
Full text available: pdf(1.34 MB) Additional Information: full citation, abstract, ret

Dynamic superscalar processors execute multiple instructions out-of-order by l window. The number of physical registers within the processor has a direct imp instructions require a new physical register at dispatch. A large multi-ported re parallelism (ILP), but may have a detrimental effect on clock speed, especially

66 Untrusted hosts and confidentiality: secure program partitioning
Steve Zdancewic, Lantian Zheng, Nathaniel Nystrom, Andrew C. Myers
October 2001 ACM SIGOPS Operating Systems Review , Proceedings of the eightee principles, Volume 35 Issue 5
Full text available: pdf(1.36 MB) Additional Information: full citation, abstract, re

This paper presents secure program partitioning, a language-based technique t computation in distributed systems containing mutually untrusted hosts. Confic by annotating programs with security types that constrain information flow; th automatically to run securely on heterogeneously trusted hosts. The resulting c implement the original p ...

67 Parallel execution of prolog programs: a survey
Gopal Gupta, Enrico Pontelli, Khayri A.M. Ali, Mats Carlsson, Manuel V. Hermene
July 2001 ACM Transactions on Programming Languages and Systems (TOPLAS
Full text available: pdf(1.95 MB) Additional Information: full citation, abstract, refere

Since the early days of logic programming, researchers in the field realized the in the execution of logic programs. Their high-level nature, the presence of no among other characteristics, make logic programs interesting candidates for ot the same time, the fact that the typical applications of logic programming freq

Keywords: Automatic parallelization, constraint programming, logic programmi

68 Tools for application-oriented performance tuning
John Mellor-Crummey, Robert Fowler, David Whalley
June 2001 Proceedings of the 15th international conference on Supercompu
Full text available: pdf(397.34 KB) Additional Information: full citation, abstract, refe

Application performance tuning is a complex process that requi information and correlating it with source code to pinpoint the Existing performance tools don't adequately support this proce discuss some of the critical utility and usability issues for applic in the context of two performance tools, *MHSim* and *HPCView*,

69 ?-coral: a multigrain, multithreaded processor architecture
Mark N. Yankelevsky, Constantine D. Polychronopoulos
June 2001 Proceedings of the 15th international conference on Supercompu
Full text available: pdf(196.56 KB) Additional Information: full citation, abstract, ref

Recently popularized hardware multithreading (HMT) architectu do not provide flexible and efficient methods of thread manage The &agr;-Coral architecture is a tool for investigation of a mor management. Unlike other architectures, there are no strict rec threads, and no static partitioning of resources. &agr;-Coral pr multiprogramming an ...

Keywords: multithreaded, parallelizing compiler, processor arc

70 Integrating superscalar processor components to implement register cach
Matthew Postiff, David Greene, Steven Raasch, Trevor Mudge
June 2001 Proceedings of the 15th international conference on Supercompu
Full text available: pdf(146.37 KB) Additional Information: full citation, abstract, referen

A large logical register file is important to allow effective compi windowed space of registers to allow fast function calls. Unfortu be slow, particularly in the context of a wide-issue processor w register file, and many read and write ports. Previous work has used to address this problem. This paper proposes a new regist

71 External memory algorithms and data structures: dealing with **massive da**
Jeffrey Scott Vitter
June 2001 ACM Computing Surveys (CSUR), Volume 33 Issue 2
Full text available: pdf(828.46 KB) Additional Information: full citation, abstract, referen

Data sets in large applications are often too massive to fit completely inside th input/output communication (or I/O) between fast internal memory and slower performance bottleneck. In this article we survey the state of the art in the des algorithms and data structures, where the goal is to exploit locality in order to

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hashing, exte multidimensional access methods, multilevel memory, online, out-of-core, secc

72 Characterizing the memory behavior of Java workloads: a structured view
Yefim Shuf, Mauricio J. Serrano, Manish Gupta, Jaswinder Pal Singh
June 2001 ACM SIGMETRICS Performance Evaluation Review , Proceedings of the on Measurement and modeling of computer systems, Volume 29 Issue
Full text available: pdf(1.55 MB) Additional Information: full citation, abstract, re

This paper studies the memory behavior of important Java workloads used in t based on instrumentation of both application and library code in a state-of-the about these workloads to help guide systems' design. We begin by characterizi benchmarks, such as information on the breakup of heap accesses among diffe to fields and met ...

73 Measuring experimental error in microprocessor simulation
Rajagopalan Desikan, Doug Burger, Stephen W. Keckler
May 2001 ACM SIGSOFT Software Engineering Notes , Proceedings of the 2001 sy software reuse in context, Volume 26 Issue 3
Full text available: pdf(1.03 MB) Additional Information: full citation, referenc

74 Locality vs. criticality

Roy Dz-ching Ju, Alvin R. Lebeck, Chris Wilkerson

May 2001 ACM SIGARCH Computer Architecture News , Proceedings of the 28th a architecture, Volume 29 Issue 2

Full text available: pdf(960.89 KB) Additional Information: full citation, abstract, referer

Current memory hierarchies exploit locality of references to rec processor performance. Locality based schemes aim at reducin to ignore the nature of misses. This leads to a potential mis-ma and latencies realized using a traditional memory system. To b critical and non-critical. A load that needs to complete early to

75 Dead-block prediction & dead-block correlating prefetchers

An-Chow Lai, Cem Fide, Babak Falsafi

May 2001 ACM SIGARCH Computer Architecture News , Proceedings of the 28th a architecture, Volume 29 Issue 2

Full text available: pdf(972.60 KB) Additional Information: full citation, abstract, referer

Effective data prefetching requires accurate mechanisms to pre blocks to prefetch and “when” to prefetch them. Predictors (DBPs), *trace-based predictors that accurately identi cache block becomes evictable or “dead”. Predict prefetching lookahead and opportunity, and enables placing da*

76 Concurrency, latency, or system overhead: which has the largest impact performance?

Vinodh Cuppu, Bruce Jacob

May 2001 ACM SIGARCH Computer Architecture News , Proceedings of the 28th a architecture, Volume 29 Issue 2

Full text available: pdf(904.17 KB) Additional Information: full citation, abstract, referer

Given a fixed CPU architecture and a fixed DRAM timing specifi for a DRAM system organization. Parameters include the numb of each channel, burst sizes, queue sizes and organizations, tu page protocol, algorithms for assigning request priorities and s this design space, we see a wide variation in application execut

77 Cache decay: exploiting generational behavior to reduce cache leakage p
Stefanos Kaxiras, Zhigang Hu, Margaret Martonosi
May 2001 ACM SIGARCH Computer Architecture News , Proceedings of the 28th a architecture, Volume 29 Issue 2
Full text available: pdf(1.17 MB) Additional Information: full citation, abstract, reference

Power dissipation is increasingly important in CPUs ranging fro way up to high-performance processors for high-end servers. V dynamic switching power, leakage power is also beginning to b future chip generations, leakage's proportion of total chip powe

This paper examines methods for reducing leakage power within the

78 A Web Odyssey: from Codd to XML
Victor Vianu
May 2001 Proceedings of the twentieth ACM SIGMOD-SIGACT-SIGART symposium
Full text available: pdf(282.10 KB) Additional Information: full citation, references, citings

79 Execution-based prediction using speculative slices
Craig Zilles, Gurindar Sohi
May 2001 ACM SIGARCH Computer Architecture News , Proceedings of the 28th a architecture, Volume 29 Issue 2
Full text available: pdf(1.03 MB) Additional Information: full citation, abstract, reference

A relatively small set of static instructions has significant levera These problem instructions contribute a disproportionate numb mispredictions because their behavior cannot be accurately ant branch prediction mechanisms.

The behavior of many problem instructions can be predicted by exe speculative slice. If a speculative slice is exec ...

80 A design framework to efficiently explore energy-delay tradeoffs
William Fornaciari, Donatella Sciuto, Cristina Silvano, Vittorio Zaccaria
April 2001 Proceedings of the ninth international symposium on Hardware/soft
Full text available: pdf(511.37 KB) Additional Information: full citation, abstract, ref

Comprehensive exploration of the design space parameters at t evaluate architectural tradeoffs accounting for both energy and we propose a system-level design methodology for the efficient architecture from the energy-delay combined perspective. The configuration of the memory hierarchy without performing the space. The target ...